



Image designed by SPC Kelly Lowery, USA

DECISION SUPPORT FOR BEST PRACTICES: LESSONS LEARNED ON BRIDGING THE GAP BETWEEN RESEARCH AND APPLIED PRACTICE

Raimund L. Feldmann, Forrest Shull, and Michele A. Shaw

Today, everyone is looking at best practices for developing a system or making the right choice in acquiring system components. If the right best practices are applied, they help to avoid common problems and improve quality, cost, or both. However, finding and selecting an appropriate best practice is not always an easy endeavor. In most cases guidance, based on sound experience, is missing; often the best practice is too new, still under study, or the existing experiences do not fit the user's context. This article reports on a program that tries to bridge the gap between rigorous empirical research and practical needs for guiding practitioners in selecting appropriate best practices.

Many program managers would agree that using time-tested “Best Practices” can help to avoid common problems and increase the quality of a system, reduce development cost, or both. For instance, in a short survey at the 2004 Conference on the Acquisition of Software-Intensive Systems, 48 senior systems and software managers supported the use of Best Practices. However, the same survey indicated that it is hard to find such Best Practices. The survey identified the following reasons for this problem:

- Best practices often do not exist (i.e., they have not been publicly documented),

- People do not know of a certain best practice, or
- Best practices are not easily accessible (i.e., there is no central place to look for best practices).

The last point matches a more general study by the Delphi Group in which more than 65 percent of the interviewees agreed that finding the right information to do their job is difficult (Delphi, 2002).

Further research conducted by the U.S. Department of Defense (DoD) concluded that barriers for the adoption of best practices included:

- the lack of selection criteria among practices within cost-constrained programs,
- the lack of confidence in the value of such practices by the program offices, and
- the inability to relate practices to the risks and issues programs were facing.

In summary, recognizing good practices and disseminating them to the workforce seems to be a key issue. To address these issues the DoD Acquisition Best Practices Clearinghouse (BPCh) program, sponsored by several offices of the DoD (DS, ARA, National Information Infrastructure [NII], and Defense Procurement & Acquisition Policy [DPAP]), was initiated in 2003 (Dangle, Dwinnell, Hickok & Turner, 2005).

The Fraunhofer Center for Experimental Software Engineering, Maryland (FC-MD) was chosen to develop the initial “proof of concept” for a system to document, evaluate, and disseminate Best Practices. In collaboration with other organizations within the DoD and industry (including Northrop Grumman IT, the Computer Sciences Corporation [CSC], and the Systems and Software Consortium [SSCI]), a prototype system has been built and piloted. It is currently operated and hosted by the Defense Acquisition University (DAU).

THE VISION FOR APPLYING BEST PRACTICES

The DoD vision for the BPCh initiative is to provide more than just a list of Best Practices. It is to provide an integrated set of processes, tools, and resources which will enable information seekers to identify emerging or well-proven practices that have been implemented and proven effective. Practices in the BPCh serve as an information resource to individuals looking for ideas on how to improve quality and become more effective in their job. Clearly, the vision of the BPCh is not to create another “data cemetery,” but to develop an information-sharing network around the BPCh repository which will foster relationships between individuals within DoD and also partnerships between DoD and industry leaders. The following types of questions illustrate usage examples:

- “I just heard about accelerated life testing. Where can I find out if it’s useful or just hype?”
- “They’ve just shortened my testing schedule by 30 percent. Are there any practices that can help me better handle that kind of schedule compression?”
- “I want to add inspections to my quality process. Is it worth the cost and if so, what’s a good first step? Is there someone I can contact in case of any difficulties?”
- “I’ve taken over an acquisition program just before Critical Design Review (CDR). What practices should I look for in my contractors?”
- “I’m in charge of defining a training course as part of the continuing education program for quality improvements. What are state-of-the-art or emerging practices that should be addressed?”

The BPCh has been designed with the understanding that a single practice can never be a “silver bullet” for each and every project/program. This is because some practices may only be useful or beneficial in certain contexts while failing to produce the desired results in others. For example, practices that are absolutely necessary for large, mission critical projects may be too heavyweight for rapid prototyping or Web application development. Practices that work well when the development team is located in the same room may not always scale well when the team is distributed across the country.

Clearly, there exists no one “best” answer. Practices that are best for one user might not be best for the next. Therefore, the BPCh tool responds to user queries with a list of practices rated by how well they fit the project characteristics of the user making the query. The presented selection is compiled using the experience other users have had implementing the practice in a similar context. High-quality evidence about a practice is collected and reported with any necessary caveats, so that information seekers have a sound basis for making up their own minds given their needs.

APPLYING TECHNOLOGY TO DELIVER BEST PRACTICES

To develop the BPCh tool, we applied FC-MD’s EMPEROR approach (Experience Management Portal using Empirical Results as Organizational Resources). This approach makes use of all kinds of available evidential data from research and industry, analyzes and packages it, and disseminates it through a Web-based Experience Base.

The EMPEROR is based on the *experience factory* approach, developed by Basili, Caldiera, and Rombach (1994), which has been successfully employed to facilitate organizational learning at NASA (Basili, et al., 1995), DaimlerChrysler (Schneider & Schwinn, 2001), and elsewhere in North America, Europe, and Australia (Koennecker, Jeffery, & Low, 2000; Mendonca, Seaman, Basili, & Kim,

2001). An experience factory provides a way to analyze results based on practical experience, and package what is learned into an Experience Base for new users of the organization to find and apply.

Since the users of the BPCh come from a wide variety of organizations and programs, any Experience Base will have difficulties in addressing all user needs. To mitigate this problem, EMPEROR is required to: (a) provide transparency to users, so that they can understand the analysis process and the sources of experience and make up their own minds; (b) rate the “trustability” of each of the used sources, so that users can judge the degree of confidence they have in the information provided; and (c) provide a completeness and maturity indicator of the practice information taken as a whole, that is, to perform a self-rating based on how much and what quality evidence can be offered.

DATA STRUCTURE OF A BPCH PRACTICE

These sections describe how these requirements are implemented in the case of the BPCh. In the BPCh, each practice has one associated Practice Record, containing information about the practice and what is available in the Clearinghouse, and zero to many Evidence Profiles, each of which contains a summary of a single organization’s experience using the practice.

A Practice Record consists of:

1. A Practice Detail block, which contains information such as the practice name, a short description, and the completeness and maturity indicator for the experience package.
2. A Practice Summary block, which synthesizes all available evidence data and describes possible application contexts for the practice based on a set of characterizing attributes. This part of the practice record thereby allows different users (i.e., organizations) to make use of the practice.

An Evidence Profile contains an example or report of some type of program that has used this practice, how they applied it, and what results were obtained. Each Evidence Profile contains the same set of context and result fields as the Practice Summary block, except that the information recorded in each field will describe only what has been observed in the given context of the particular piece of evidence. In addition, the data structure of an Evidence Profile contains a field for documenting its classification of the trustability.

TRUSTABILITY OF A SINGLE SOURCE OF EVIDENCE

A 20-point scale rates the trustability of each Evidence Profile. A rating of 1 indicates an anecdotal or informal experience; a rating of 20 indicates that the results of applying the practice are rigorously measured and substantiated. Points are based on the following four dimensions:

- how the practice was applied, ranging from a single pilot study to use on multiple real projects;
- how the results were measured, ranging from an educated guess to a rigorous measurement program;
- how the evidence was reported, ranging from an informal anecdote to a peer-reviewed publication; and
- who reported the evidence, ranging from a second-hand report to someone directly involved on the team.

More information on the rating scale can be found on the BPCh page of the Acquisition Community Connection of DAU (<https://acc.dau.mil/bpch>).

MATURITY OF A PRACTICE RECORD

A 4-point scale is used to rate each Practice Record to quickly inform the user of how much, and what type of, information is known about the practice. As required by EMPEROR, this scale focuses on the quality of the overall accumulated information that is available for a practice (i.e., the synthesized and packaged information in the Practice Record). Based on the available information we describe the practice maturity as:

- No status assigned/Initial entry: A new Practice Record is initially entered into the BPCh when it is nominated by our experts and/or user communities. Typically at this time, only some of the fields in the Practice Detail block are filled in and no Evidence Profiles are available.
- Bronze status/Awareness raised: As soon as any evidence becomes available (i.e., an Evidence Profile has been linked to the Practice Record), the status is set to Bronze Level. For users, the Bronze Level status indicates that the practice has been nominated by our experts and user communities, and received a preliminary check for applicability.
- Silver status/Evaluation performed: When a sufficient set of Evidence Profiles is available, the BPCh experts will fill in the Practice Summary block and the status is set to Silver Level. For users, the Silver Level status indicates that the practice has been selected as promising enough to commission experts in the area to summarize key information. Users can see at a glance what they should know.
- Gold status/Continuously maintained: When the summary has been further evaluated (i.e., vetted) by experts from industry, academia, and government, the status is set

to Gold Level. For users, the Gold Level status indicates that the practice has been through a rigorous analysis by a committee of experts in the practice itself as well as by user representatives. Information on Gold Level practices contains the best and widest-ranging experiences we can find.

CONTENT STATUS OF THE BPCh

We have been piloting BPCh processes and tools by seeding initial content. At this point the BPCh contains 51 practices at all levels of maturity. Practices that have progressed to Gold Level are those, like inspection/technical review, which have a long history of published industrial experience.

Many practices of interest in the area of systems and software acquisition have few documented sources of evidence or experience. Therefore, we are testing different processes for eliciting information from the workforce.

Based on the recommendations of our User Advisory Group, the following types of practices are currently our top-priority areas for additional content:

- Earned Value Management,
- Risk Management,
- Information Assurance, and
- Spiral Development Process.

We hope that visitors to the BPCh tool will try out the offered features for providing short stories about their own experience with practices in these (or any other) areas. We encourage you to provide feedback as to whether you agree or disagree with the existing experiences that have been entered, or thoughts on our BPCh tool in general.

LESSONS LEARNED

Based on our experience with the BPCh program and other knowledge-management projects, we can formulate some observations which make useful rules of thumb for good practices to build such systems. The BPCh program has been organized along three parallel (but interconnected) tracks, which reflects our first lesson learned.

LESSON 1: PROCEED IN MULTIPLE DIRECTIONS SIMULTANEOUSLY

Progress in building a knowledge repository needs to proceed in multiple dimensions simultaneously: content collection, tool development, and outreach.

Although there is often a temptation to view these as tasks that can be done sequentially (e.g., first the tool will be built, then populated, and then it will be advertised to users), we have found this to be an overly simplistic view that diminishes the chance of project success. Constructing the tool prior to collecting actual content and getting users' feedback almost ensures that important user needs will be discovered late and will require much more effort to implement. Populating the content without getting user feedback leads to a high likelihood that the content will not really address user needs. More importantly, content needs to come from the user community, if the repository is to have a long-term life. We have found that for the research team to generate substantial amounts of content is a time-consuming way of recreating what many users already have at their finger tips. Finally, engaging in outreach and building excitement in the community of potential users runs the risk of all prototyping efforts: When told how anything is possible in the final system, users often come up with many wish list features that are not really linked to their everyday needs. Moreover, users often get frustrated with the slow pace of progress when the system actually has to be implemented, and lose interest before the system is fielded.

Constructing the tool prior to collecting actual content and getting users' feedback almost ensures that important user needs will be discovered late and will require much more effort to implement.

To avoid these problems, we have adopted an incremental approach, with content and tool development going on simultaneously and outreach activities to the user community (such as booths at major conferences, or specific User Advisory Group meetings) planned at major milestones. Although this sometimes stretches resources a bit thin, we feel this approach has enabled us to engage periodically with the user community, show them progress since the last iteration, and get feedback on ever more mature versions of the system, with an initial body of content.

LESSON 2: MAINTAIN A CONTINUOUS STREAM OF FUNDING

Because of the interconnected nature of all the tasks listed above, having a stable funding stream is crucial. Requiring the team to take a hiatus from the project after a release is delivered leads to lost opportunities for user involvement (users find it hard to match their schedule to the development team's), leads to new content ideas that miss getting followed up on, may result in the loss of expertise if experienced personnel resources are in transition to other projects during the hiatus, increases the

personnel learning curve encountered at restarts, and may result in flagging interest in the user community since momentum generated during outreach is lost.

LESSON 3: RECOGNIZE THE RELATIVE MERITS OF CONTENT

Our most important lesson learned is a direct implication of the BPCh vision: There is no such thing as a “Best Practice.” Or, to say it more diplomatically: No practice will be “best” for every project. Practices that are absolutely necessary for large, mission-critical projects may be too heavyweight for rapid prototyping or Web application development.

The implications of this lesson are many. Perhaps the most important is related to the tone of the recommendations that users find: Rather than arguing as an expert that readers *should* be following a given practice, or else they are doing something wrong, practices should be recommended to readers on the basis that projects of certain type(s) have found it useful. That is, rather than presenting a foregone conclusion to users, the system should aim at respecting users’ intelligence enough to enable them to draw their own conclusion, providing sufficient evidence as necessary for those decisions to be sound ones.

LESSON 4: UNDERSTAND THE LIFE CYCLE OF BEST PRACTICES

Practices (and practice information) are not static and have a real life cycle. Major paradigm shifts in the software development world can have an impact on which practices are recommended. The practices that seemed to be good fits for most projects, when a waterfall life cycle was the most common approach to software development, are not all equally applicable at the current time, when iterative, spiral, and even agile approaches are probably more representative of the state-of-the-art practice.

Our recommendations regarding a structured life cycle for practice information are:

1. A knowledge repository needs to be continually evolving by accepting information on topics of interest and making it available to users as soon as possible. While some quality checking is necessary to make sure that incorrect, misleading, or incomplete information is disseminated outward, it is better to get information to users as it comes in, than to wait and try to create something perfect. Users should be able to see a timestamp on all information so that they can see if the experiences related are fresh and up to date or come from years ago.
2. However, the desire to get information out quickly should not interfere with the need for validation activities that provide higher confidence in the information. These additional levels of maturity should be noted, to give users more confidence in the information they find, but should not be used as a *precondition* for displaying content.

3. Content needs to be retired *when appropriate*. Practices may have a natural lifespan, since the acquisition and development worlds continue to evolve and change on their own. Practices that were good 10 years ago may not be appropriate given today's constraints or technologies. To avoid users finding obsolete information in the repository, reports need to be generated periodically of which practices have received no updates or new experiences in the longest time.

LESSON 5: APPLY AGILE STRATEGIES AND PROTOTYPING

To create the front end of the BPCh tool, which helps users find candidate practices, explore possibilities, and get more information on practices of real interest, we have found that prototyping and agile strategies are extremely valuable for developing knowledge-management systems.

***Speak to the users in their language.
Do not expect them to learn yours.***

Precisely because of the need for parallel activities in different tracks, and the number of stakeholders involved (tool developers, content gathering team, end user representatives, sponsor representatives), an agile approach is extremely valuable. The implementation of the prototype BPCh tool was carried on in two-week increments, at the end of which a releasable version was always available. At the end of each two-week period, a demonstration and planning meeting was held with as many of the stakeholders as could be present. This approach was necessary to help us coordinate and prioritize the evolving expectations of the users as well as the necessary changes that were suggested by the content development team, based on what they were finding.

As part of this meeting we learned the following lesson:

LESSON 6: USE APPROPRIATE LANGUAGE

Speak to the users in their language. Do not expect them to learn yours. We realized early on that having the greatest possible content in the BPCh repository would not be of much help if the users cannot find it. To address this we needed to provide multiple paths to the information, so that users could select the path that made the most sense to them. Some specific lessons learned here included:

1. Organize around common tasks. The best way to reach users is to organize the contents of the repository according to everyday activities that the user performs. This helps users see the repository less as an additional activity that they need to

make time for, and more as a value-added to the activities that already consume their time. In the case of BPCh, we added several such perspectives (i.e., indexes to the content) based around activities of importance to different segments of the user community (e.g., addressing CMMI practice areas, constructing a systems engineering strategy, and referencing back to common guidebooks).

2. Push as well as pull information. Rather than always expecting users to take time to come to browse the BPCh tool, information can be “pushed” outward to the user on a periodic basis. For example, the user could select some practices of special interest, and when new experiences come in related to these practices a notification is sent via e-mail.
3. Match users to practices based on context similarity. Since no practice will be “best” for every project, it is important to match users to practices using context characteristics. This provides the users with a pick list of practices that may be useful in their particular situation. In addition, it may alert the user to practices that they might not have known about previously. For example, if the user selects a few context variables that describe his/her context, then practices can be prioritized and displayed according to whether they have associated evidence provided by users with similar context information. This is a way of indicating that, even if the practice does not answer a specific search query, users like the current one have found this practice useful and it may be something the user should know.

LESSON 7: DEMONSTRATE PRACTICAL EXAMPLES TO INTENDED USER

To engage in effective outreach activities, aimed at building up an interested and active community of users of the BPCh, we find the following lesson of relevance: You can not show initial users an empty depository.

In line with the idea that building a tool like the BPCh needs to proceed on three tracks in parallel (front-end, content, and outreach) is the lesson that populating the content cannot come after the repository is built. Showing users a fancy front-end without an initial set of real content may get their interest for a short time period, but is not an effective way of building an active user community. Users need to see a small but representative set of content which they can respond to and start generating ideas for the next content or tool release.

LESSON 8: UPDATE CONTENT AND FUNCTIONALITY CONTINUOUSLY

To keep interest engaged, when users do check back to the site they need to see that updates have been made since last time. Content needs to be continuously updated and refreshed to stay abreast of trends.

If users ever become convinced that the repository does not get updated on a regular basis, this often spells the end of their involvement. Rather, they need to be motivated to come back often enough to find new things and hopefully, as they progress, be motivated to submit responses and ideas of their own showing emerging

trends and keeping the content relevant. Thus, user involvement tends to build more user involvement. As users become interested enough to post comments or send new ideas to the repository, other users will continue to be interested to show up to see which comments have been added since the last time and possibly find something of interest to their current situation—and more likely to find something applicable.

One way we have experimented with—to reinforce this concept—is to list on the front page of the BPCh tool the most recently added practices and highlight ones that have been promoted to various maturity levels (Bronze, Silver, or Gold). Thus, one of the first things users see is an indicator of how much progress has occurred since their last visit.

CONCLUSIONS

This article has presented some of the lessons learned with the BPCh program, which aims to document practices and quickly disseminate them to the users. The BPCh, which is based on the EMPEROR approach, makes use of a two-dimensional rating scale. These scales provide users with a quick overview of the trustability and maturity of the stored practice records. The scales allow users to understand and to draw their own conclusions based on a set of evidence from different contexts, from research studies as well as industrial experiences, and using measures at different levels of rigor. Practitioners can rely on this information without reading in detail through the different evidence sources, unless they are interested in the very detailed level of information.

In addition, ways to collect user feedback and trigger discussions are offered to allow a vivid and growing user community. While initial feedback regarding the BPCh tool has been positive (Turner & Shull, 2005), we are continuing to improve the BPCh program and its associated tool through ongoing research, advisory groups, and user community feedback. We are interested in addressing such questions as: “How much extra effort to certify evidence sets and summaries as correct is worthwhile to users?” or “Are there subsets or types of evidence that users will find especially worthwhile?”

We invite you to take a look at our BPCh tool, available at <http://bpch.dau.mil>. We appreciate all feedback, whether it be submitted through the tool or directly to the authors’ e-mail.

ACKNOWLEDGMENTS

This research was supported with funding from the U.S. Department of Defense (DoD), the Office of the Secretary of Defense (OSD), and the Defense Acquisition University (DAU). We wish to thank the members of the BPCh team, from DAU, FC-MD, CSC, and SSCI, for the many productive discussions that have improved this work.



Mr. Raimund L. Feldmann is the technical lead for Knowledge and Experience Management at the Fraunhofer Center for Experimental Software Engineering, MD (FC-MD). Before he joined FC-MD in 2004, Raimund participated in several technology transfer projects in Germany and was also involved in the development of the Virtual Software Engineering Competence Center (VSEK) portal, funded by the Department of Education and Research (bmb+f) of the German Federal Government, to offer up-to-date Software Engineering knowledge to subject matter experts.

(E-mail address: rfeldmann@fc-md.umd.edu)



Mrs. Michele A. Shaw is a Scientist at the Fraunhofer Center for Experimental Software Engineering. Michele supports clients implementing process improvement, measurement, and experience factory concepts. She has over 25 years of experience in Information Technology including software and service development, project management, quality assurance, client care and subcontractor management. Ms. Shaw holds a BS in Business from University of Baltimore and a masters in applied behavioral science from Johns Hopkins University.

(E-mail address: mshaw@fc-md-umd.edu)



Dr. Forrest Shull is a senior scientist at the Fraunhofer Center for Experimental Software Engineering, MD (FC-MD). He is project manager and member of technical staff for projects with clients that have included Fujitsu, Motorola, NASA, and the U.S. Department of Defense. He has also been lead researcher on grants from the National Science Foundation, Department of Energy, Air Force Research Labs, and NASA's Office of Safety and Mission Assurance.

(E-mail address: fshull@fc-md.umd.edu)

AUTHOR BIOGRAPHY

REFERENCES

- Basili, V. R., Caldiera, G., & Rombach, H. D. (1994). Experience factory. In J. J. Marciniak (Ed.), *Encyclopedia of Software Engineering* (Vol. 1, pp. 469–476). New York: John Wiley & Sons, Inc.
- Basili, V., Zelkowitz, M., McGarry, F., Page, J., Waligora, S., & Pajerski, R. (1995). SEL's software process improvement program. *IEEE Software*, 12(6), 83–87.
- Dangle, K., Dwinnell, L., Hickok, J., & Turner, R. (2005, May). Introducing the Department of Defense acquisition best practices clearinghouse. *CrossTalk*, 18(5), 4–5.
- Defense Acquisition University. Retrieved from <http://bpch.dau.mil>
- Delphi White Paper. (2002). Taxonomy & Content Classification—Market Milestone Report. Boston, MA: Delphi Group.
- Koennecker, A., Jeffery, R., & Low, G. (2000, April). Implementing an experience factory based on existing organizational knowledge. In *Proceedings of the 2000 Australian Software Engineering Conference* (pp. 28–29), Canberra, ACT, Australia.
- Mendonca, M., Seaman, C., Basili, V. R., & Kim, Y. M. (2001, June). A prototype experience management system for a software consulting organization. In *Proceedings of the 13th International Conference on Software Engineering and Knowledge Engineering (SEKE)*. Ottawa, Canada.
- Schneider, K., & Schwinn, T. (2001, June). Maturing experience base concepts at DaimlerChrysler. *Software Process—Improvement and Practice*, 6(2), 85–96.
- Turner, R., & Shull, F. (2005, November). An empirical approach to best practice identification and selection: The U.S. Department of Defense acquisition best practices clearinghouse. In *Proceedings of the 4th International Symposium on Empirical Software Engineering (ISESE 2005)*(pp. 133–140), Noosa Heads, Australia.